# Accurate modeling of plasma acceleration with arbitrary order pseudo-spectral particle-in-cell methods

S. Jalas,[1, a)] I. Dornmair,[1] R. Lehe,[2] H. Vincenti,[2] J.-L. Vay,[2] M. Kirchen,[1] and A. R. Maier[1]
[1)] *Center for Free-Electron Laser Science & Department of Physics, University of Hamburg, 22761 Hamburg, Germany*
[2)] *Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA*

Particle in Cell (PIC) simulations are a widely used tool for the investigation of both laser- and beam-driven plasma acceleration. It is a known issue that the beam quality can be artificially degraded by numerical Cherenkov radiation (NCR) resulting primarily from an incorrectly modeled dispersion relation. Pseudo-spectral solvers featuring infinite order stencils can strongly reduce NCR – or even suppress it – and are therefore well suited to correctly model the beam properties. For efficient parallelization of the PIC algorithm, however, localized solvers are inevitable. Arbitrary order pseudo-spectral methods provide this needed locality. Yet, these methods can again be prone to NCR. Here, we show that acceptably low solver orders are sufficient to correctly model the physics of interest, while allowing for efficient parallelization.

PACS numbers: 02.70.-c, 52.65.Rr, 29.27.-a

arXiv:1611.05712v1 [physics.acc-ph] 17 Nov 2016

## INTRODUCTION

Plasma accelerators provide high accelerating gradients and promise very compact next-generation particle accelerators as drivers for brilliant light sources[1,2] and high energy physics[3]. Using laser or beam drivers, the acceleration of electron beams by multiple GeV over few-cm distances has been demonstrated in experiments[4–7]. The improvement of beam quality, especially in terms of energy spread, remains crucial for applications. Due to the non-linear nature of the beam-plasma interaction, Particle-in-Cell (PIC) simulations are a vital tool to model the beam injection and transport in the plasma. PIC algorithms self-consistently solve the interaction of spatially discretized electromagnetic fields with charged particles defined in a continuous phase space. However, these codes can suffer from numerical Cherenkov radiation (NCR)[8] that is caused primarily by incorrect modeling of the electromagnetic dispersion relation. This, for example, has been found to artificially decrease the beam quality in terms of transverse beam emittance[9] for finite-difference time domain (FDTD) solvers, such as the commonly used Yee scheme. Efforts to circumvent this issue include the modification of the FDTD dispersion relation[9,10] and digital filtering[11,12]. In contrast to FDTD solvers, pseudo-spectral solvers, which solve Maxwell's equations in spectral space, strongly reduce NCR – and sometimes even suppress it.

In practice, PIC simulation can easily demand thousands of hours of computation time. Therefore, implementations that are efficient and scalable for parallel production are unavoidable. The parallelization of PIC codes is typically done by decomposing the simulation volume into domains which are evaluated by individual processes. For FDTD algorithms efficient scaling to several thousand parallel processes is possible. Pseudo-spectral solvers, however, act globally on the grid due to the needed Fourier transform, which limits efficient parallelization. An approach to overcome this limitation are *arbitrary order* pseudo-spectral solvers[13–15], which only locally affect the fields when solving Maxwell's equations. On the downside, this — again — introduces spurious numerical dispersion and can make a simulation prone to NCR.

In the following, we investigate numerical Cherenkov radiation and its implications to the physics of a simulation for arbitrary order pseudo-spectral solvers using the spectral, quasi-cylindrical PIC algorithm implemented in the code FBPIC[16].

The paper is structured as follows. We will first summarize the concept of pseudo-spectral solvers, with an emphasis on the *arbitrary order PSATD*[13], followed by a discussion of the main characteristics of NCR. In the last part we numerically analyze the mitigation of NCR in the arbitrary order PSATD scheme. As a figure of merit for the physical accuracy of the solver we use the effect of NCR on the beam emittance and the slice energy spread.

## FINITE ORDER PSEUDO-SPECTRAL SOLVERS

A widely used method in PIC codes are FDTD solvers. These algorithms use finite difference operators to approximate derivatives in time and space when numerically solving Maxwell's equations, which introduces a spurious numerical dispersion. This approximation can be done with arbitrary precision by increasing the used stencil order, i.e., taking contributions from more surrounding grid nodes into account. Yet, this has the drawback of also increasing the computational demand of the calculation. Further, common FDTD algorithms represent the electromagnetic fields on a staggered grid, where the electric and magnetic fields are not defined at the same points in time and space. This is known to cause errors when modeling the direct interaction between an

[a)] Electronic mail: soeren.jalas@desy.de

electron beam and a co-propagating laser[17].

The use of finite differences in space can be avoided by transforming Maxwell's equations into spectral space. In that case the differential operator $\partial_z$ is given by a multiplication with $ik_z$, which numerically can be evaluated without approximations. The resulting solver class is referred to as *pseudo-spectral time domain* (PSTD)[18] and it describes the limit of an FDTD solver with an infinite stencil order[14]. Yet, this solver still applies finite difference approximations in time and thereby, like FDTD solvers, its timestep is limited by a Courant condition. It is subject to spurious numerical dispersion, but unlike FDTD schemes the erroneous dispersion caused by the PSTD scheme results in modes traveling faster than the speed of light[19].

Under the assumption of constant currents $\boldsymbol{J}$ over one timestep $\Delta t$, this constraint can be overcome by integrating the Maxwell equations analytically in time. With a precise representation of derivatives in time and space, the resulting *pseudo-spectral analytic time domain* (PSATD)[20] scheme is not subject to a Courant condition and is free of spurious numerical dispersion. Furthermore, it is currently the only solver that can be combined with the Galilean scheme[21,22] that is intrinsically stable when modeling relativisticly drifting plasmas with uniform velocity.

The PSATD field propagation equations in Cartesian coordinates are

$$\begin{aligned}\tilde{\boldsymbol{E}}^{n+1} &= C\tilde{\boldsymbol{E}}^{n} + iS\hat{\boldsymbol{k}}\times\tilde{\boldsymbol{B}}^{n} - \frac{S}{ck}\tilde{\boldsymbol{J}}^{n+1/2} \\ &\quad + i\frac{\hat{\boldsymbol{k}}}{k}\left[\left(\frac{S}{ck\Delta t}-1\right)\tilde{\rho}^{n+1} + \left(C-\frac{S}{ck\Delta t}\right)\tilde{\rho}^{n}\right] \qquad (1a)\\ \tilde{\boldsymbol{B}}^{n+1} &= C\tilde{\boldsymbol{B}}^{n} - iS\hat{\boldsymbol{k}}\times\tilde{\boldsymbol{E}}^{n} + i\frac{1-C}{ck}\hat{\boldsymbol{k}}\times\tilde{\boldsymbol{J}}^{n+1/2}, \qquad (1b)\end{aligned}$$

where $\boldsymbol{k}$ is the wave vector of length $k \equiv \sqrt{k_x^2+k_y^2+k_z^2}$ with $\hat{\boldsymbol{k}} = \boldsymbol{k}/k$ and $\tilde{F}$ is the Fourier transform of a field $F$. Further, $C = \cos(kc\Delta t)$ and $S = \sin(kc\Delta t)$. With pseudo-spectral methods the fields are naturally centered in space and time, eliminating a source of interpolation errors due to the usual staggering of the electromagnetic fields in the popular Yee FDTD solver. The equations of the PSTD scheme can be retrieved from those of the PSATD scheme, by writing the PSATD equations in an equivalent staggered form[19] and performing a first-order Taylor expansion in $\Delta t$.

In the following we will use the PSATD scheme in a quasi-cylindrical geometry as derived in Ref.[16]. In this geometry the fields are decomposed into azimuthal modes $m$, which can then be represented on separate 2D grids[23]. In practice, as the contributions from most of the modes go to zero for near-symmetrical systems, only a few of these have to be taken into account. Consequently, this geometry allows to model near-symmetrical problems at a drastically reduced computational effort compared to full 3D simulations. In the quasi-cylindrical PSATD scheme the real-space components $F_r$, $F_\theta$ and $F_z$ of a field $\boldsymbol{F}$ are connected to their spectral components $\tilde{F}_{+,m}$, $\tilde{F}_{-,m}$ and $\tilde{F}_z$ by the combination of a Hankel transform (along $r$) and a Fourier transform (along $z$). The equations of the quasi-cylindrical PSATD can be derived from the Cartesian PSATD scheme by replacing the differential operators in Eqs. (1a) and (1b) by[22]

$$\tilde{\boldsymbol{F}} = i\boldsymbol{k}\tilde{S} \rightarrow \begin{pmatrix}\tilde{F}_{+,m}\\ \tilde{F}_{-,m}\\ \tilde{F}_{z,m}\end{pmatrix} = \begin{pmatrix}-k_\perp\tilde{S}_m/2\\ k_\perp\tilde{S}_m/2\\ ik_z\tilde{S}_m\end{pmatrix} \qquad (2a)$$

$$\tilde{\boldsymbol{F}} = i\boldsymbol{k}\times\tilde{\boldsymbol{V}} \rightarrow \begin{pmatrix}\tilde{F}_{+,m}\\ \tilde{F}_{-,m}\\ \tilde{F}_{z,m}\end{pmatrix} = \begin{pmatrix}k_z\tilde{V}_{+,m} - ik_\perp\tilde{V}_{z,m}/2\\ -k_z\tilde{V}_{-,m} - ik_\perp\tilde{V}_{z,m}/2\\ ik_\perp\tilde{V}_{+,m} + ik_\perp\tilde{V}_{-,m},\end{pmatrix} \qquad (2b)$$

where $\tilde{S}$ is any scalar field and $\tilde{\boldsymbol{V}}$ any vector field. This representation of the PSATD scheme preserves the advantages of the Cartesian PSATD scheme, i.e. an ideal dispersion relation, a centered representation of the fields and no Courant condition for the timestep, and combines them with the reduced computational demand associated with a quasi-cylindrical geometry.

The propagation of the fields in spectral space in Eqs. (1a) and (1b) corresponds to a global operation in spatial space, where information is transferred over the entire simulation grid. However, for domain decomposition a locality of the solver operations is necessary. In order to achieve an arbitrary order local stencil the solver equations are modified.

The formulations of the modified pseudo-spectral scheme can be derived from the expression of a finite difference with arbitrary order. For a centered scheme the discrete derivative with stencil order $2\,n$ of a function $F$ at a discrete point $p$ on the spatial grid is[24]

$$\begin{aligned}(D_zF)_p &= \sum_{j=1}^{n}\alpha_{n,j}\frac{F_{p+j}-F_{p-j}}{2\,j\,\Delta z} \qquad (3)\\ \alpha_{n,j} &= (-1)^{j+1}\frac{2\,(n!)^2}{(n-j)!\,(n+j)!}.\end{aligned}$$

To get a spectral method with a finite stencil derivative consider the spectral representation of this operator

$$(D_zF)_p = \sum_{k_z} i[k_z]\tilde{F}_{k_z}e^{ik_zp\Delta z}. \qquad (4)$$

Here $[k_z]$ is the modified wavenumber given as

$$[k_z] = \sum_{j=1}^{n}\alpha_{n,j}\frac{\sin(k_zj\Delta z)}{j\,\Delta z}, \qquad (5)$$

with the grid spacing $\Delta z$ and $\tilde{F}$ the spectral form of $F$. This means that the equivalent to a finite stencil in real

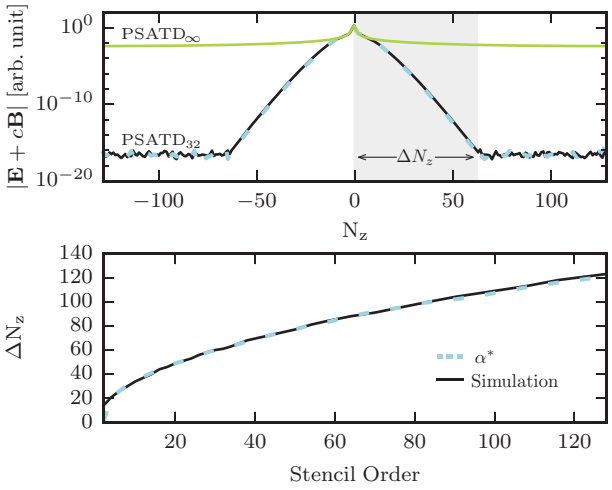


FIG. 1. Expansion of an electromagnetic unit pulse after one PIC cycle with modified $k_z$ (top). The field pushed with the unmodified PSATD method (green line) shows non-negligible field contributions along the entire simulation grid, while the fields modeled with order 32 PSATD (black line) decrease by 16 orders of magnitude over $\Delta N_z$ cells and reach the machine precision level. The number of cells over which the fields decrease (bottom) depends on the stencil order. The dashed lines show the stencil coefficients derived from the modified PSATD equations.

space can be achieved by using a modified PSTD algorithm, whereby $k_z$, $k_x$, $k_y$ are replaced by $[k_z]$, $[k_x]$, $[k_y]$ in the discretized Maxwell equations in spectral space. Because this modified PSTD algorithm is mathematically equivalent to an FDTD scheme (with high-order spatial derivatives), it has the same locality.

Similarly to the PSTD algorithm, in order to obtain a PSATD scheme with improved locality, we follow identical steps. More precisely, the discretized Maxwell equations Eqs. (1a) and (1b) are modified by replacing $k_z$, $k_x$, $k_y$ by $[k_z]$, $[k_x]$, $[k_y]$ – including in the expression of $k \equiv \sqrt{k_x^2 + k_y^2 + k_z^2}$ which in turn impacts the expressions of the coefficients $C$ and $S$. Consequently, unlike the modified PSTD algorithm, the modified PSATD scheme is not strictly local. In order to evaluate its locality, we get the corresponding stencil coefficients on the spatial grid, by applying an inverse Fourier transform to the operators used to advance the fields in Eqs. (1a) and (1b), i.e. $[\hat{\boldsymbol{k}}]S$ and $C$.

Next, we will use this scheme in quasi-cylindrical geometry with the PIC framework FBPIC to verify the locality of the solver. In FBPIC only $k_z$ is modified, because the Fourier transform is only used along the $z$-axis. Therefore, the corresponding solver stencil remains infinite in radial direction and locality is only gained in longitudinal direction. To examine this, a $\delta$-peak in E and B is initialized on axis in the center of the simulation box with a grid resolution of $\Delta r = \Delta z = 1\,\mu$m. It is then propagated for one timestep $\Delta t = \Delta z/c$. Fig. 1 (top) shows the expansion of the pulse for a stencil order of 32 compared to the standard PSATD. When using the unmodified PSATD (PSATD$_\infty$) scheme the simulation shows field contributions along the entire axis, whereas the fields modeled with finite order PSATD quickly decrease in amplitude and reach the machine precision level. The expansion of the unit pulse strictly follows the stencil coefficients $\alpha*$ derived from the modified PSATD operations given by

$$\alpha^* = \sqrt{|\mathcal{F}^{-1}(C)^2| + |\mathcal{F}^{-1}(Sk_\perp)^2| + |\mathcal{F}^{-1}(S[k_z])^2|}, \quad (6)$$

where $\mathcal{F}^{-1}$ denotes an inverse Fourier transform along the $z$-axis. Here these coefficients are calculated numerically and instead of an analytical Fourier transform an FFT was used. The coefficients are evaluated for $k_\perp = 0.5\,\mu\text{m}^{-1}$. The stencil coefficients as well as the unit pulse reach machine precision related noise after decreasing by 16 orders of magnitude over a range of $\Delta N_z$ cells. In order to show how the stencil extent $\Delta N_z$ scales with the arbitrary order PSATD, the same simulation is done with varied stencil order. Fig. 1 (bottom) shows that the stencil extent behaves non-linearly to the applied stencil order. While for high orders the extent is similar to the stencil order, the reach of the solver tends to be bigger than the order in low order cases.

In a domain decomposition scheme the fields would only need to be exchanged between neighboring processes in guard regions with the size of $\Delta N_z$ cells, as this is the maximum distance information can travel within one timestep. These guard cells hold copies of the fields in the neighboring domains. At each timestep the fields are evaluated independently in each domain and the information is only exchanged in these cells. When significantly less than $\Delta N_z$ guard cells are used, truncation of the stencil can lead to spurious fields[14].

Note that for parallelization also a localized current correction or a charge conserving current deposition is needed. For that an exact current deposition in k-space for Cartesian coordinates[19] and a local FFT based current correction on staggered grids[15] have been shown.

Compared to FDTD algorithms, arbitrary order pseudo-spectral methods allow to freely choose the stencil order, and thereby the precision of the field solver, without increasing the computational demands of the algorithm. Nevertheless, this scheme imposes approximations on the integration of the electromagnetic fields which also result in spurious numerical dispersion. The implications of this will be discussed in the following.

## NUMERICAL CHERENKOV RADIATION

In general, the phase velocity $v_\phi$ of an electromagnetic mode with wave vector $\boldsymbol{k}$ and angular frequency $\omega$ is $v_\phi^2 = \omega^2/\boldsymbol{k}^2$. The allowed electromagnetic modes in vacuum are described by the dispersion relation

$$\omega^2 = c^2\boldsymbol{k}^2. \quad (7)$$

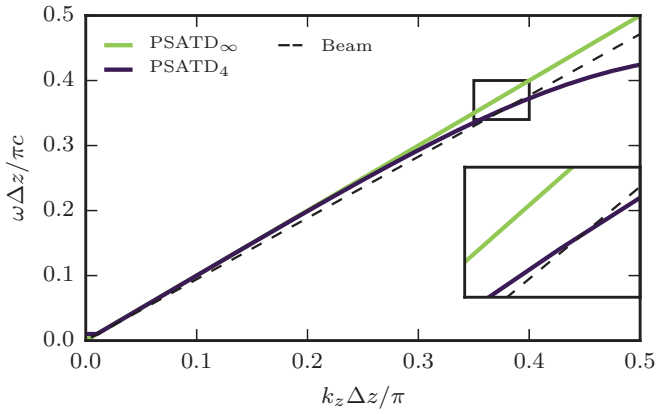


FIG. 2. Dispersion relation for the arbitrary order PSATD solver (solid lines) on the $k_z$-Axis ($k_r = 0$). Finite stencils introduce a spurious dispersion relation, which leads to a reduced phase velocity $v_{\phi,z} = \omega/k_z$ of modes with large $k_z$ and allows unphysical intersections (box) with the mode corresponding to the velocity of a relativistic particle (dashed line), here $\gamma_{\text{beam}} = 3$. Particle beams can then travel with the same velocity as electromagnetic modes, which leads to coherent emission of spurious radiation, referred to as NCR.

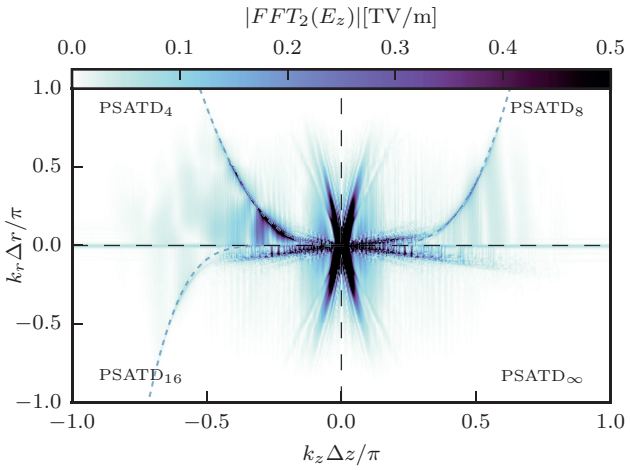


FIG. 3. Spatial Fourier transforms of the accelerating field for simulations with different PSATD stencil orders. As the Fourier transform is point symmetric around the origin only one quadrant from each simulation is shown above. The simulations with finite order show resonances around the modes with a phase velocity matching the velocity of the electron beam (dashed lines) given by Eq. (9). The magnitude of NCR decreases for increasing stencil order. The NCR-free infinite order PSATD simulation is shown for reference.

When modifying $\boldsymbol{k}$ in longitudinal direction the dispersion relation is

$$\omega^2 = c^2 \left(k_r^2 + [k_z]^2\right). \tag{8}$$

Here, a spurious numerical dispersion is introduced, where the phase velocity in vacuum decreases for modes with increasing $k_z$. Fig. 2 shows the distorted dispersion relation caused by the arbitrary order PSATD solver compared to the accurate case (Eq. (7)) given by the unmodified, infinite order PSATD method. Due to the distortion of the dispersion relation, relativistic particles can travel with the same velocity as some electromagnetic modes. This can lead to coherent emission of spurious radiation referred to as numerical Cherenkov radiation. Consequently, the affected modes obey

$$v_{\phi,z} = v_{\text{beam}}, \tag{9}$$

where $v_{\text{beam}}$ is the velocity of the charged particles in the simulation. Using Eq. (8) this can be rewritten as

$$v_{\text{beam}} = \frac{c\sqrt{k_r^2 + [k_z]^2}}{k_z}. \tag{10}$$

This condition describes the modes, where the spurious electromagnetic dispersion relation intersects with the beam mode corresponding to the velocity of a charged particle. This effect is illustrated in Fig. 3. The dashed lines indicate the modes fulfilling Eq. (9) and thus correspond to the intersection points of the dispersion relation and the beam mode as depicted in Fig. 2. These intersection points denote the modes where NCR is excited. They are overlaid with the spatial Fourier transform of the longitudinal electric field from FBPIC simulations, with parameters as discussed in the next section.

The quadrants of Fig. 3 show results from simulations done with different stencil orders. They all show excited modes that can be attributed to the physics in the simulation of a beam driven plasma accelerator, e.g. the plasma wave or betatron radiation. However, compared to the unmodified PSATD scheme simulation, additional excited modes are visible around the modes prone to NCR. The excitations are especially strong for low solver orders, whereas they vanish for PSATD$_\infty$ (i.e. no dashed line). Next, we will show the influence of these unphysical modes on the properties of an electron beam.

## ACCURATE MODELING OF PLASMA ACCELERATORS IN FINITE ORDER PSATD

To show that already low stencil orders are sufficient to mitigate errors due to NCR in the context of plasma-based acceleration, PIC simulations for typical beam-driven plasma accelerator parameters are presented. For the same physical parameters, the stencil order is varied in the spectral, quasi-cylindrical code FBPIC.

The propagation of cold, i.e. zero emittance, monoenergetic driver and witness electron beams through a plasma of density $1 \times 10^{17}$ cm$^{-3}$ is simulated. A non-linear wakefield is driven by a 1 GeV Gaussian beam with a charge of 180 pC, 8.4 µm rms length and 5 µm rms width, compare Fig. 4(a). The 25 MeV Gaussian witness bunch of 50 pC charge, 2 µm rms length and 1 µm rms width is positioned in the accelerating region at the back of the bubble.

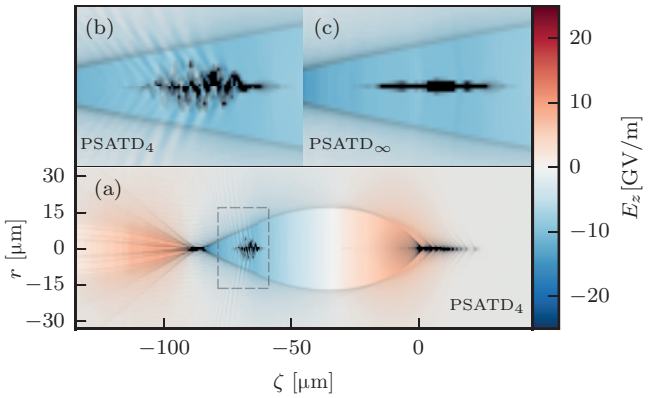


FIG. 4. (a) Simulation setup featuring a driver and witness beam after propagating 35 mm. Colors correspond to the accelerating field and the charge density is given in gray (arb. units), $\zeta = z - ct$ are the co-moving coordinates with $\zeta = 0$ at the center of the driver beam. The short PSATD$_4$ stencil leads to distortions of the longitudinal field around the witness beam (b) that are not present for an infinite order stencil (c).

The simulation box has a resolution of 5 cells/µm longitudinally and 1.5 cells/µm transversally, with 16 particles per cell distributed as $2 \times 2 \times 4$ in longitudinal, radial and azimuthal direction, respectively. A total of 900 $\times$ 300 cells in two azimuthal modes is used. While the pseudo-spectral solver of FBPIC is technically not limited by a CFL condition, a timestep of $\Delta t = \Delta z/c$ was used. The NCR-free simulation with an infinite order stencil is used as a reference in the following.

Fig. 4 and 5 both show simulation results after a propagation through 35 mm of plasma. In Fig. 4 (b) and (c) the accelerating field $E_z$ and the charge density $\rho$ in the vicinity of the witness bunch can be seen. In the case of a low stencil order strong distortions are visible in $E_z$, which affect the witness beam shape. This corresponds to an increase of beam emittance, as it has been reported for 3D FDTD codes[9].

We find that also the longitudinal phase space is affected, as shown in Fig. 5, which compares the phase spaces modeled with several stencil orders. The projected energy spread increases during propagation through the plasma due to the slope of $E_z$ in combination with beam loading, so that an s-shape like structure is imprinted on the phase space also with PSATD$_\infty$. However, for decreasing stencil order, NCR causes an increasingly violent high frequency oscillation on the longitudinal phase space, leading to an unphysical growth of both slice and projected energy spread. It can be observed that the perturbation is especially pronounced behind the beam center, where the current density is highest, while the head remains mostly unaffected.

Fig. 6 (top) shows the evolution of the slice energy spread during the propagation through the plasma for different solver orders. For finite order stencils the energy spread grows non-linearly with $z$. While PSATD$_4$ and PSATD$_8$ still deviate significantly from the Cherenkov free infinite order case, the impact of NCR on the beam in the PSATD$_{16}$ simulation is barely visible even after a long propagation distance. The dashed lines correspond to simulations done with the FDTD Yee scheme in 3D Cartesian and quasi-cylindrical geometry. These are performed with the PIC code WARP[25] using the same parameters and resolution as FBPIC. With the Yee scheme the growth rate of spurious slice energy spread is even greater than for the low order PSATD case. Note that the WARP framework also features a PSATD$_\infty$ (and arbitrary order) solver, which, if used for this problem, would also produce the correct physical results as the FBPIC PSATD$_\infty$ algorithm.

Fig. 6 (bottom) shows the slice energy spread and the transverse beam emittance after a propagation length of 35 mm depending on the stencil order. The beam's energy spread is even more sensitive to potential NCR than the beam emittance. Yet, both quantities quickly converge to the PSATD$_\infty$ result for increasing stencil order. The transverse and longitudinal phase space are modeled correctly for stencil orders higher than 12 and 32, respectively.

We expect NCR to become more significant with higher charge, therefore a more precise dispersion relation, i.e. a higher stencil order, is potentially necessary in such a scenario. We find that for a witness beam charge of 100 pC, though not explicitly shown here, a stencil order of 64 is sufficient to mitigate the unphysical growth of slice energy spread over the same propagation distance of 35 mm. Further the growth of NCR is highly reduced by increasing the longitudinal resolution of the simulation grid. When doubling the longitudinal resolution in the presented case, PSATD$_8$ is already sufficient to preserve the physics like in the PSATD$_\infty$ reference case.

## CONCLUSION

In conclusion, NCR stemming from erroneous modeling of the dispersion relation leads to beam phase space degradation, which manifests itself in a spurious growth of slice energy spread and emittance. This effect is present in both the FDTD Yee scheme and finite order PSATD solvers. Errors from NCR rapidly decrease for higher stencil orders, eventually showing the same results as the NCR-free infinite order PSATD solver.

In the case discussed here, which is inspired by typical beam driven plasma acceleration parameters, a stencil order of 32 effectively suppresses spurious beam quality degradation. As a conservative example a long propagation distance of 35 mm and high witness beam charge of 50 pC was chosen. Since NCR grows in time and also scales with the beam charge, it can be expected that this stencil order is also sufficient for more relaxed parameter sets than shown here.

In practice, this means that with around 60 guard cells between neighboring domains, parallelization is possible

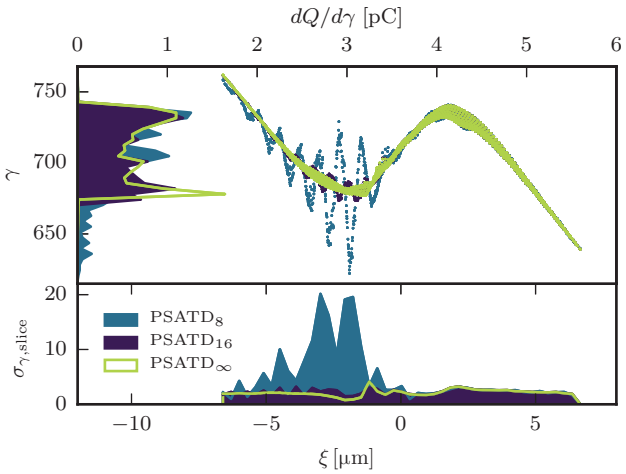


FIG. 5. Longitudinal phase space of the witness beam after propagation through 35 mm of plasma. $\xi$ is the internal bunch coordinate with $\xi = 0$ at the center of the witness beam. For decreasing stencil order it is increasingly spoiled by a high frequency oscillation. This causes a growth of projected and slice energy spread (slice length of 0.3 µm here).

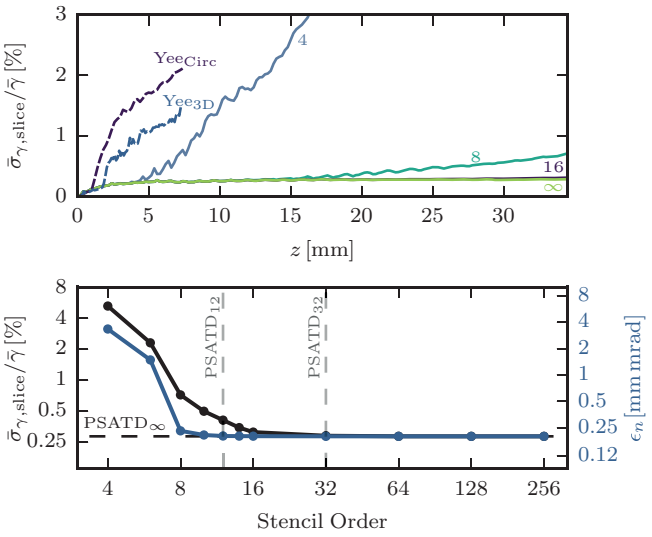


FIG. 6. Top: Growth of the witness beam's slice energy spread during the propagation through 35 mm of plasma. The slice energy spread is averaged over all slices within ± one standard deviation around the bunch center. The solid and dashed lines show data obtained with the PSATD and Yee scheme, respectively. The PSATD$_{16}$ and PSATD$_\infty$ lines almost overlay each other. Bottom: For increasing stencil order the final normalized emittance and slice energy spread after 35 mm quickly converge to the infinite order case.

while retaining the physics of the considered problem. Analogous to the frequently used Yee scheme, growth of NCR in the finite order stencil PSATD scheme can also be reduced by increasing the resolution of the simulation grid. However, this is an inefficient solution especially in schemes where the timestep is limited by a CFL-condition and hence directly linked to the grid resolution. Yet, this means that in simulations with high spatial resolution the effects of NCR are less severe. This, for example, would be the case for laser plasma acceleration, where a high resolution is required to resolve the laser wavelength. Therefore, in these cases even lower order stencils than suggested here are sufficient to model the physics without artifacts from NCR.

The arbitrary order PSATD scheme preserves the benefits of pseudo-spectral solvers, e.g. the analytic integration in time or a centered definition of the electric and magnetic fields. It further allows to independently adjust the precision of the electromagnetic dispersion relation and the resolution of the simulation. This way the needed spatial resolution is governed by the physical problem and not by the mitigation of NCR.

## ACKNOWLEDGMENTS

We gratefully acknowledge the computing time provided on the supercomputer JURECA under project HHH20. Work at LBNL was funded by the Director, Office of Science, Office of High Energy Physics, U.S. Department of Energy under Contract No. DE-AC02-05CH11231, including the Laboratory Directed Research and Development (LDRD) funding from Berkeley Lab.

[1] A. R. Maier, A. Meseck, S. Reiche, C. B. Schroeder, T. Seggebrock, and F. Grüner, Phys. Rev. X **2**, 031019 (2012).
[2] Z. Huang, Y. Ding, and C. B. Schroeder, Phys. Rev. Lett. **109**, 204801 (2012).
[3] C. B. Schroeder, E. Esarey, C. G. R. Geddes, C. Benedetti, and W. P. Leemans, Phys. Rev. ST Accel. Beams **13**, 101301 (2010).
[4] W. P. Leemans, B. Nagler, A. J. Gonsalves, C. Toth, K. Nakamura, C. G. R. Geddes, E. Esarey, C. B. Schroeder, and S. M. Hooker, Nat Phys **2**, 696 (2006).
[5] W. P. Leemans, A. J. Gonsalves, H.-S. Mao, K. Nakamura, C. Benedetti, C. B. Schroeder, C. Tóth, J. Daniels, D. E. Mittelberger, S. S. Bulanov, J.-L. Vay, C. G. R. Geddes, and E. Esarey, Phys. Rev. Lett. **113**, 245002 (2014).
[6] M. Litos, E. Adli, W. An, C. I. Clarke, C. E. Clayton, S. Corde, J. P. Delahaye, R. J. England, A. S. Fisher, J. Frederico, S. Gessner, S. Z. Green, M. J. Hogan, C. Joshi, W. Lu, K. A. Marsh, W. B. Mori, P. Muggli, N. Vafaei-Najafabadi, D. Walz, G. White, Z. Wu, V. Yakimenko, and G. Yocky, Nature **515**, 92 (2014).
[7] X. Wang, R. Zgadzaj, N. Fazel, Z. Li, S. A. Yi, X. Zhang, W. Henderson, Y. Y. Chang, R. Korzekwa, H. E. Tsai, C. H. Pai, H. Quevedo, G. Dyer, E. Gaul, M. Martinez, A. C. Bernstein, T. Borger, M. Spinks, M. Donovan, V. Khudik, G. Shvets, T. Ditmire, and M. C. Downer, Nat Commun **4** (2013).
[8] B. B. Godfrey, Journal of Computational Physics **15**, 504 (1974).
[9] R. Lehe, A. Lifschitz, C. Thaury, V. Malka, and X. Davoine, Phys. Rev. ST Accel. Beams **16**, 021301 (2013).
[10] B. M. Cowan, D. L. Bruhwiler, J. R. Cary, E. Cormier-Michel, and C. G. R. Geddes, Phys. Rev. ST Accel. Beams **16**, 041303 (2013).
[11] A. D. Greenwood, K. L. Cartwright, J. W. Luginsland, and E. A. Baca, Journal of Computational Physics **201**, 665 (2004).
[12] J.-L. Vay, C. Geddes, E. Cormier-Michel, and D. Grote, Journal of Computational Physics **230**, 5908 (2011).

[13] J.-L. Vay and A. Arefiev, AIP Conference Proceedings **1777**, 030002 (2016), http://dx.doi.org/10.1063/1.4965596.
[14] H. Vincenti and J.-L. Vay, Computer Physics Communications **200**, 147 (2016).
[15] F. Li, P. Yu, X. Xu, F. Fiuza, V. K. Decyk, T. Dalichaouch, A. Davidson, A. Tableman, W. An, F. S. Tsung, R. A. Fonseca, W. Lu, and W. B. Mori, ArXiv e-prints (2016), arXiv:1605.01496 [physics.comp-ph].
[16] R. Lehe, M. Kirchen, I. A. Andriyash, B. B. Godfrey, and J.-L. Vay, Computer Physics Communications **203**, 66 (2016).
[17] R. Lehe, C. Thaury, E. Guillaume, A. Lifschitz, and V. Malka, Phys. Rev. ST Accel. Beams **17**, 121301 (2014).
[18] Q. H. Liu, Microwave and Optical Technology Letters **15**, 158 (1997).
[19] J.-L. Vay, I. Haber, and B. B. Godfrey, Journal of Computational Physics **243**, 260 (2013).
[20] I. Haber, R. Lee, H. Klein, and J. Boris, *Proc. Sixth Conf. on Num. Sim. Plasmas, Berkeley, CA* (1973) pp. 46 – 48.
[21] M. Kirchen, R. Lehe, B. B. Godfrey, I. Dornmair, S. Jalas, K. Peters, J.-L. Vay, and A. R. Maier, Physics of Plasmas **23** (2016), http://dx.doi.org/10.1063/1.4964770.
[22] R. Lehe, M. Kirchen, B. B. Godfrey, A. R. Maier, and J.-L. Vay, Phys. Rev. E **94**, 053305 (2016).
[23] A. Lifschitz, X. Davoine, E. Lefebvre, J. Faure, C. Rechatin, and V. Malka, Journal of Computational Physics **228**, 1803 (2009).
[24] B. Fornberg, Mathematics of computation **51**, 699 (1988).
[25] J.-L. Vay, D. P. Grote, R. H. Cohen, and A. Friedman, Computational Science & Discovery **5**, 014019 (2012).